arXiv:1604.06133v1 [cs.LG] 20 Apr 2016

Embedded all relevant feature selection with Random Ferns

Miron Bartosz Kursa
Interdisciplinary Centre for Mathematical and Computational Modelling
University of Warsaw
Pawińskiego 5A, 02-106 Warsaw, Poland

April 22, 2016

Abstract

Many machine learning methods can produce variable importance scores expressing the usability of each feature in context of the produced model; those scores on their own are yet not sufficient to generate feature selection, especially when an all relevant selection is required. Although there are wrapper methods aiming to solve this problem, they introduce a substantial increase in the required computational effort.

In this paper I investigate an idea of incorporating all relevant selection within the training process by producing importance for implicitly generated shadows, attributes irrelevant by design. I propose and evaluate such a method in context of random ferns classifier. Experiment results confirm the effectiveness of such approach, although show that fully stochastic nature of random ferns limits its applicability either to small dimensions or as a part of a broader feature selection procedure.

1 Introduction

The crucial part of any machine learning application is to find a good representation of the observed data; its ability to express the contained information in a way that is learnable for the utilised modelling method is often the most important component of the final performance. On the other hand the robustness of this process is also critical, as it is very easy to introduce overfitting this way; either by leaking information which model should not be given or by amplifying false, random associations which are only present in training data due to its finite size.

The are also cases when the original data is already in a 'tabular' form, i.e. as a series of independently measured features, and it is strongly desired to retain this structure because of known links between features and certain physical aspects of the investigated phenomenon. Then, the only possible form of representation altering is *feature selection*, reducing the original set of attributes to a relevant, sometimes also non-redundant subset. Such selection may be less effective in terms of learnability, yet becomes an useful result on its own by highlighting important aspects of the problem.

Most prominent examples here are data sets obtained via high throughput biological experiments: they can simultaneously capture the activity of thousands or even millions of agents representing a substantial fraction of a full state of a given system, though only a handful is expected to be connected to the investigated state or behaviour. Finding important yet previously unknown agents in such case can lead (through targeted studies) to a discovery of novel mechanisms, consequently be more important than the original task of building a black-box model.

Fundamentally, feature selection methods can be divided into two classes, *minimal optimal* and *all relevant* [NPnBT07] — methods of the first group attempt to find a smallest subset of features on which certain model achieves optimal performance; second group collects methods which attempt to remove features irrelevant to the problem, consequently retaining those features which may be useful for modelling. The first aim is straightforward to implement as an accuracy optimisation problem; unfortunately, such selection is susceptible to contain attributes forming false associations, on the other hand missing true, interesting interactions due to their apparent redundancy. Consequently, only all relevant selection is applicable for deciphering mechanisms behind the analysed data; also it leads to a better robustness of the selection, especially in a $p \gg n$ setting.

There is also a technical taxonomy due to how the selection is coupled with modelling [SILn07]: there are *filters*, algorithms which are independent from the modelling method, *embedded methods* which integrate selection and modelling into a single algorithm, finally *wrappers* which relay on some modelling method, but only as an oracle of an efficiency of a given subset of features or feature importance. Obviously all filters have to be all relevant methods, yet they often only analyse univariate or very simple interactions between attributes and decision, leading to poor results [SAP08]. There are wrappers belonging to each of fundamental classes, though the universal problem of those methods is their computational demand, as they usually require building hundreds of models. Embedded methods are generally much faster, though they are usually either minimal optimal or only offer variable

importance scores (VIM) or ranking rather than strict selection.

All relevant wrappers [TBT06, HTWG08, KJR10] generally relay on external feature importance provider, yet try to establish a VIM threshold separating relevant and irrelevant features. While it is mostly impossible to analytically obtain the distribution of VIM for irrelevant variables, aforementioned wrappers use a kind of permutation test approach to estimate it — utilising either permutation of the decision attribute or features irrelevant by design (called *shadows* or *contrasts*) injected within original ones. Such an approach obviously requires many repetitions to stabilise the approximation; some methods also try to progressively eliminate irrelevant features, as both stability and the quality of VIM often decrease with the dimensionality of a set.

The aim of this work is to propose a methodology to implement such all relevant selection in an embedded method setting, bringing together selection quality and computational efficiency. In particular, I will closely follow the heuristics behind the Boruta method [KR10], as it is generic and proved effective in a demanding assessment [Kur14b]. While it can be in principle generalised for any method producing variable importance or ranking, here I will introduce and evaluate it in context of random ferns, an efficient, stochastic ensemble classifier.

2 Boruta

Boruta [KJR10, KR10] is an all relevant feature selection wrapper, originally developed for Random Forest [Bre01], but in a present form capable of using arbitrary classifier that outputs VIM. The algorithm works in a following way. First, VIM is calculated on a modified data set extended with explicitly generated shadows, i.e. nonsense features created from the original ones by permuting the order of values within them (thus wiping out information but preserving distribution). Then, all original variables which importance was larger than the maximal importance of a shadow are assigned a *hit*.

This procedure is repeated (with shadows re-generated each time), while the proportion of hits for each feature is observed for being significantly lower or higher than half; in the first case the feature is claimed rejected and removed from further consideration, in the latter claimed confirmed and added to final result. The algorithm stops either when the status of all original features is decided or when a previously given maximal number of iterations is exhausted, in which case some variables may be left undecided.

3 Random Ferns

Random ferns are classifiers introduced in [OFL07] and named as such in [OCLF08]. They were developed to serve as an computationally efficient alternative to Random Forests in demanding computational vision tasks [BZM07, OGIB09]. I have previously shown, however, that with certain modifications they can perform well in a generic machine learning context [Kur14a], as well as produce variable importance measure (VIM) of a similar quality as Random Forest, though in a much shorter time [Kur14b]. Here I will briefly present the method in the aforementioned generalised version, called *rFerns*.

Let's assume a set of N training objects $\mathcal{T} \subset \mathcal{X}$, where $\mathcal{X}$ is composed of M attributes, i.e., $\mathcal{X} = \prod_{j=1}^{M} \mathcal{X}_j$, where $\mathcal{X}_j$ is a domain of the j-th attribute. Each object $\mathbf{x} \in \mathcal{T}$ is assigned $Y(\mathbf{x})$, one of C disjoint classes.

rFerns is an ensemble of K ferns $\mathbf{F}^k(\mathcal{T}') : \mathcal{X} \to \mathbb{R}^C$ returning a vector of class scores indicating confidence how a given object fits within each class. The ensemble is built using bagging, thus k-th fern is built on a separate random multiset of training objects called *bag*, $\mathcal{B}_k := B(\mathcal{T}, k)$, where $B(\mathcal{X}, \lambda)$ denotes sampling with replacement the same number of elements as in the input set, and λ functions as a random seed of sampling procedure, i.e., samples with different λ would be different and statistically independent. The prediction of the ensemble is the class which gets maximal sum of scores over all ferns, i.e.,

$$\hat{Y}(\mathbf{x}) := \arg\max_y \sum_k F_y^k(\mathbf{x}; \mathcal{B}_k). \tag{1}$$

The base of $\mathbf{F}$ is a *trunk* function $T : \mathcal{X} \to 1..2^D$ classifying given object into its corresponding *leaf*, which is then associated with a vector of classes' scores estimated using $\mathcal{T}$. D is a hyperparameter of the classifier and is called *fern depth*. Trunk is defined by a sequence of D split attribute indexes $j_i \in 1..M$ and subsets $\Xi_i \subset \mathcal{X}_{j_i}$, so that an element $\mathbf{x} = (x_1, x_2, \ldots, x_M) \in \mathcal{X}$ belongs to a leaf

$$T(\mathbf{x}) := \sum_{i=0}^{D} 2^i \cdot I(x_{j_i} \in \Xi_i), \tag{2}$$

where I is an indicator function. As rFerns is a highly stochastic method, trunks are generated randomly: D split attributes are drawn uniformly with replacement from $1..M$, and subsets Ξ_i are either random subsets of $\mathcal{X}_{j_i}$ (for finite, unordered domains) or generated from some random threshold value θ_i, i.e. $\Xi_i = \{x \in \mathcal{X}_{j_i} : x \geq \theta_i\}$.

On the other hand, scores are based on the distribution of training objects' classes over leaves, adjusted for the class imbalance and with add-one smoothing. For convenience, let's first define

$$\mathcal{L}_{\mathcal{B},T}(\mathbf{x}) = \{\xi \in \mathcal{B} : T(\xi) = T(\mathbf{x})\}\,, \tag{3}$$

a subset of bag objects in the same leaf as $\mathbf{x}$ and

$$\mathcal{Y}_{\mathcal{B}}(y) = \{\xi \in \mathcal{B} : Y(\xi) = y\}\,, \tag{4}$$

a subset of bag objects of a class y. Full fern is then defined as

$$\exp(F_y(\mathbf{x};\mathcal{B})) := \frac{1 + \#(\mathcal{L}_{\mathcal{B}}(\mathbf{x}) \cap \mathcal{Y}_{\mathcal{B}}(y))}{C + \#\mathcal{L}_{\mathcal{B}}(\mathbf{x})} \cdot \frac{C + \#\mathcal{B}}{1 + \#\mathcal{Y}_{\mathcal{B}}(y)}, \tag{5}$$

where # denotes multiset cardinality. Note that the score expresses the deviation of class proportion within leaf with respect to what is expected from random assignment; under-represented classes have negative scores while over-represented positive. For balanced classes and $\mathcal{L}_{\mathcal{B}}(\mathbf{x}) = \varnothing$, the score is $\mathbf{0}$.

4 Shadow Importance

Similarly to Random Forest, the regular rFerns VIM is defined using the *out-of-bag* or OOB objects, i.e., those which are not in the bag of a given fern; I denote this set here as $\mathcal{B}^* := \mathcal{T} - \mathcal{B}$. rFerns VIM for an attribute a is defined as

$$I_a = \sum_{k \in \mathcal{A}(a)} \frac{1}{\#\mathcal{A}(a) \cdot \#\mathcal{B}_k^*} \left(\sum_{\xi \in \mathcal{B}_k^*} F^k_{Y(\xi)}(\xi;\mathcal{B}_k) - \sum_{\xi^\diamond \in \mathcal{B}_{k,a}^{*\diamond}} F^k_{Y(\xi^\diamond)}(\xi^\diamond;\mathcal{B}_k) \right), \tag{6}$$

where $\mathcal{A}(a)$ is a set of ferns that incorporate feature a, and $\mathcal{B}_{k,a}^{*\diamond}$ is $\mathcal{B}_k^*$ in which values within attribute a has been randomly shuffled, thus decoupled from other attributes' values and object classes, consequently making a irrelevant. The involved permutation is different for every (k, a) pair.

In the Boruta method, distribution of VIM of irrelevant attributes is estimated by calculating VIM on the training set augmented with shadows, randomly permuted copies of all true features. This method is straightforward but inefficient — augmented set occupies twice as much memory, also the search space of potential interactions becomes significantly larger. To this end the shadow attributes for rFerns shadow VIM are created implicitly

during the importance calculation loop, utilising fern trunks built only on actual features. This yields *shadow importance*, which is defined as

$$\begin{aligned} J_a = &\sum_{k \in \mathcal{A}(a)} \frac{1}{\#\mathcal{A}(a) \cdot \#\mathcal{B}_k^*} \sum_{\xi \in \mathcal{B}_k^*} F_{Y(\xi)}^k(\xi; B(\mathcal{T}_a^\diamond, k)) + \\ &- \sum_{k \in \mathcal{A}(a)} \frac{1}{\#\mathcal{A}(a) \cdot \#\mathcal{B}_k^*} \sum_{\xi^\diamond \in \mathcal{B}_{k,a}^{*\diamond}} F_{Y(\xi^\diamond)}^k(\xi^\diamond; B(\mathcal{T}_a^\diamond, k)), \end{aligned} \tag{7}$$

where $\mathcal{T}_a^\diamond$ is $\mathcal{T}$ in which values of an attribute a have been randomly shuffled. Note that while $\mathcal{B}_{k,a}^{*\diamond}$ is different for each (k, a) pair, $\mathcal{T}_a^\diamond$ is the same for all ferns; this is crucial because implicit shadows are required to mimic the behaviour of original attributes, which obviously do not change due to being used by different ferns.

To implement feature selection with shadow importance, each original attribute importance is compared with the maximum shadow importance, i.e. the selected subset of features is

$$\mathcal{S} = \{a \in 1..M : I_a > \max_{a'} J_{a'}\}; \tag{8}$$

this approach also follows the heuristics behind the Boruta method.

5 Numerical Experiments

The proposed feature selection system is assessed on a series of 6 synthetic problems built based on an established benchmark datasets. All of them belong to $p \gg n$ class, and contain *a priori* known set of relevant attributes; in case of three of them this set is empty, as they were derived through total randomisation of a benchmark data. Precisely:

- Iri — a derivative of the [Fis36] iris data, enlarged by adding 4996 irrelevant features generated by shuffling original ones.
- Iri2 — similar to Iri, but with 1000 irrelevant features only.
- Mad — the Madelon dataset from the Nips 2003 feature selection challenge [GGBHD05], a 5-dimentional xor problem extended with 15 random linear combinations of 5 main attributes (also considered relevant in this work) and 480 irrelevant features containing random values.

- RND — a derivative of the [GST+99] microarray data, made nonsense by randomly shuffling values within each feature.

- RND2 — a derivative of the [SFR+02] microarray data, made nonsense in a same way as in case of RND.

- RND3 — another nonsense derivative of the Singh et al. data, however made by randomly shuffling values within decision (so that all inter-attribute relation present in the original data are retained).

For each of those datasets, a series of rFerns-based embedded feature selections has been performed: over 10 repetitions with different random seeds and over a comprehensive subspace of hyperparameters. The investigated values of the depth parameter D were 1, 3, 5, 7, 9, 10 and 12. For sake of comparability between cases for different D and sizes of the data, the ensemble size parameter K was set to value corresponding to a given average number of considerations (*scans*) of a feature, i.e. $\langle \#\mathcal{A}(a) \rangle_a$; here the scan has been performed for values in range 10–5000.

As a base for comparison, a similar analysis has been performed with Boruta using, as in its original implementation, Random Forest VIM; it was also stabilised by taking 10 repetitions and applied for 500, 5000 and 50000 trees in ensemble. This is because Random Forest has no parameter corresponding to fern depth, while the number of scans is not uniform across features due to a greedy generation of trees. To compare the results of Boruta and ferns, the average, single core running time of both algorithms was used as an indication of the profoundness of a method in a place of the average attribute scan count.

The measured accuracy of rFerns-based selections is presented on Figure 1, and compared with the Boruta result on Figure 2. One can see that in all cases except MAD the selection is very accurate; the number of false negatives always converge to 0, while the number of false positives is of an order of several attributes, which corresponds to false positive rates of and order of 10^{-4}. The quality of the selection is also robust across values of D, although deeper ferns require more attribute scans (consequently ferns in ensemble) to converge; overall the selection is most effective for an intermediate value of $D = 7$.

The Boruta method yields more accurate result only in case of the MAD set, and generates few times more false positives than rFerns selection for all sets. When considering computational time, rFerns always managed to converge withing 10 minutes, while Boruta took at minimum 1 hour and 9.4

Figure 1: The accuracy of rFerns all relevant feature selection for benchmark problems, shown as the number of false positive and false negative attribute selections. Position on the x axis corresponds to to the average number of times each attribute was considered in importance calculation (which is approximately K/D); columns correspond to various fern depths while rows to problems. Note that for RND* problems the number of false negatives is strictly 0 by design.

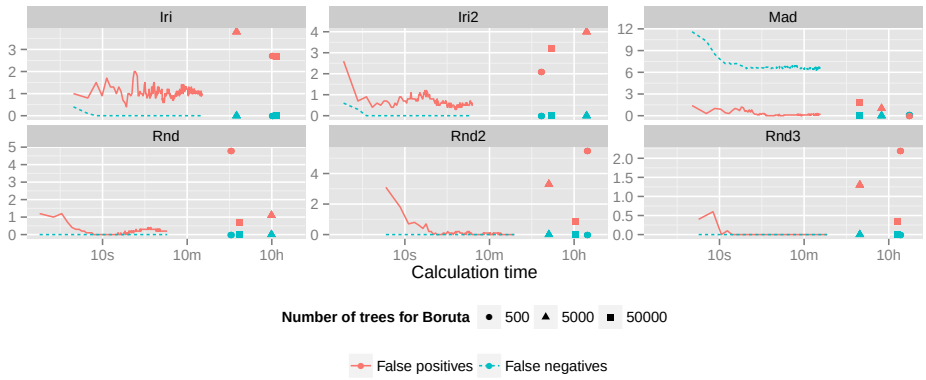

Figure 2: Data as on Figure 1, yet shown only for $D = 7$, with respect to a single core calculation time and compared with the results of a Boruta feature selection (points).

hours at average[1]. It is also worth noticing that the selection error in case of RND3 (in which all attributes are noise) is substantially smaller than in case of RND2 (in which only the decision is permuted).

Because of an inferior performance of rFerns with comparison to Boruta on the MAD data, I decided to study this problem deeper, namely how the selection accuracy depends on the number of nonsense features in the training set. To this end, the original data was spread into a family of derived problems MAD_w of a controlled complexity — namely, MAD_w contains only w of 480 irrelevant features present in the original Madelon dataset. This test was performed by applying rFerns feature selection on MAD_w for $w \in \{1, 2, 3, \ldots, 50, 55, 60, \ldots, 150, 200, 250, \ldots, 450\}$ and, similarly to previous experiments, using 10 repetitions, for $D \in \{1, 3, 5, 7, 9, 10, 12\}$ and for the average number of scans equal to 50, 100, 200, 500, 1000 and 2000.

The results of this analysis are summarised on Figure 3. One can see that the XOR problem underlying MAD data is solvable perfectly by the rFerns method for $D \geq 3$, although up to only around 100 irrelevant features ($5\times$ the number of relevant ones) in case of $D = 7$ and more than 1000 scans of each attribute. Still, for a reasonable amount of scans the method never misses more than 7 features; while those are the same attributes that $D = 1$

[1]One should note that those are single-core timings, while both Random Forest used by Boruta and rFerns are trivially parallel and scale well even up to hundreds of processing cores.

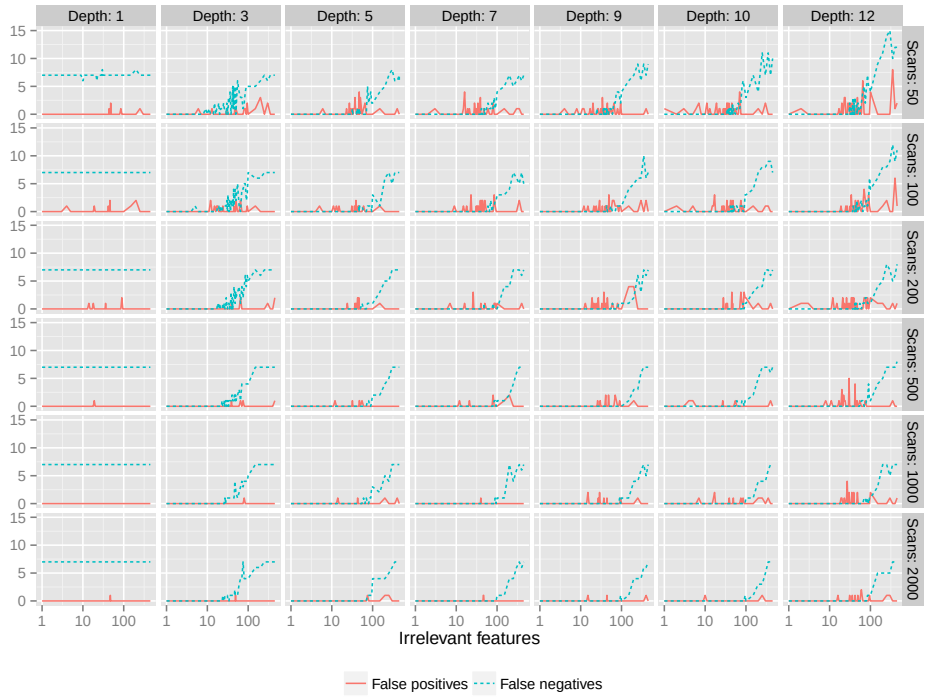

Figure 3: The results of rFerns all relevant feature selection for a family of Madelon dataset modifications made by trimming the original set of 480 irrelevant attributes to a random subset of from 1 to 480 of them (corresponding to the position on the x axis). The plot shows the number of false positive and false negative attribute selections; columns correspond to various fern depths, while rows to the average number of times each attribute was considered in importance calculation (which is approximately K/D).

rFerns always miss, it is likely that those are relevant features which can only be identified in multivariate analysis. Rare false positives seem to arise at random, with probability rising with D and the amount of noise in the system; though they can be practically eliminated by increasing the number of scans.

Those results suggests that the poor performance of rFerns selection on the full MAD problem is caused by the fact that trunks are built stochastically and not optimised, as in Random Forest utilised by Boruta, thus is not fundamental to the construction of the feature selection methodology.

6 Conclusions

The paper proposes an approach to extend basic variable importance score into an embedded, all relevant feature selection method. It is based on the construction of the Boruta wrapper method, in particular the idea of shadow attributes; they are though generated only implicitly, which leads to a better efficiency and smaller memory footprint.

The results of numerical assessment of the obtained method show that it is highly specific, robust and computationally efficient. On a range of problems it can par or even outperform Boruta, still achieving substantial speed-ups. Due to a fully stochastic nature of the underlying rFerns classifier, its capability of detecting complex multivariate interactions is though susceptible to a high dimensionality of the input set. Consequently, it seems that applying rFerns feature selection for complex problems with hundreds or more of features would require adding some amount of optimisation; for instance, trunk building may be biased to include important features more often, or the entire method wrapped in a forward selection scheme.

Acknowledgements

This work has been financed by the National Science Centre, grant 2011/01/N/ST6/07035, as well as with the support of the OCEAN — *Open Centre for Data and Data Analysis* Project, co-financed by the European Regional Development Fund under the Innovative Economy Operational Programme.

Computations were performed at the ICM, grant G48-6.

References

[Bre01] Leo Breiman. Random Forests. *Machine Learning*, 45:5–32, 2001.

[BZM07] Anna Bosch, Andrew Zisserman, and Xavier Munoz. Image Classification Using Random Forests and Ferns. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.

[Fis36] R. A. Fisher. The Use of Multiple Measurements in Taxonomic Problems. *Annals of Eugenics*, 7(2):179–188, 1936.

[GGBHD05] I. Guyon, S. Gunn, A. Ben-Hur, and G. Dror. Result analysis of the nips 2003 feature selection challenge. *Advances in Neural Information Processing Systems*, 17:545–552, 2005.

[GST+99] T. R. Golub, D. K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J. P. Mesirov, H. Coller, M. L. Loh, J. R. Downing, M. A. Caligiuri, C. D. Bloomfield, and E. S. Lander. Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring. *Science*, 286(5439):531–537, October 1999.

[HTWG08] Van Anh Huynh-Thu, Louis Wehenkel, and Pierre Geurts. Exploiting tree-based variable importances to selectively identify relevant variables. *JMLR: Workshop and Conference Proceedings*, pages 60–73, 2008.

[KJR10] Miron B Kursa, Aleksander Jankowski, and Witold R Rudnicki. Boruta – A System for Feature Selection. *Fundamenta Informaticae*, 101(4):271–285, 2010.

[KR10] Miron B. Kursa and Witold R. Rudnicki. Feature selection with the Boruta package. *Journal of Statistical Software*, 36(11):1–13, 2010.

[Kur14a] Miron Bartosz Kursa. rFerns: An Implementation of the Random Ferns Method for General-Purpose Machine Learning. *Journal of Statistical Software*, 61(10):1–13, 2014.

[Kur14b] Miron Bartosz Kursa. Robustness of Random Forest-based gene selection methods. *BMC Bioinformatics*, 15(1):8, January 2014.

[NPnBT07] R Nilsson, J.M. Peña, J. Björkegren, and J. Tegnér. Consistent feature selection for pattern recognition in polynomial time. *The Journal of Machine Learning Research*, 8:612, 2007.

[OCLF08] Mustafa Özuysal, Michael Calonder, Vincent Lepetit, and Pascal Fua. Fast Keypoint Recognition using Random Ferns. *Image Processing*, 2008.

[OFL07] Mustafa Özuysal, Pascal Fua, and Vincent Lepetit. Fast Keypoint Recognition in Ten Lines of Code. *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007.

[OGIB09] Olusegun Oshin, Andrew Gilbert, John Illingworth, and Richard Bowden. Action Recognition Using Randomised Ferns. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference*, pages 530–537. IEEE, 2009.

[SAP08] Yvan Saeys, Thomas Abeel, and Yves Van de Peer. Robust Feature Selection Using Ensemble Feature Selection Techniques. In Walter Daelemans, Bart Goethals, and Katharina Morik, editors, *Machine Learning and Knowledge Discovery in Databases*, number 5212 in Lecture Notes in Computer Science, pages 313–325. Springer Berlin Heidelberg, January 2008.

[SFR+02] Dinesh Singh, Phillip G Febbo, Kenneth Ross, Donald G Jackson, Judith Manola, Christine Ladd, Pablo Tamayo, Andrew a Renshaw, Anthony V D'Amico, Jerome P Richie, Eric S Lander, Massimo Loda, Philip W Kantoff, Todd R Golub, and William R Sellers. Gene expression correlates of clinical prostate cancer behavior. *Cancer cell*, 1(2):203–9, March 2002.

[SILn07] Yvan Saeys, Iñaki Inza, and Pedro Larrañaga. A review of feature selection techniques in bioinformatics. *Bioinformatics*, 23(19):2507–17, October 2007.

[TBT06] Eugene Tuv, Alexander Borisov, and K. Torkkola. Feature Selection Using Ensemble Based Ranking Against Artificial Contrasts. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pages 2181–2186. IEEE, 2006.